

What is claimed is:

- Sub AI
- 1 1. A method of dynamically modifying a fragmentation size cluster
2 communication parameter in a clustered computer system, the method comprising:
3 (a) deferring processing of a requested fragmentation size change until
4 receipt of an acknowledgment message for at least one unacknowledged
5 message sent by a source node to a plurality of target nodes;
6 (b) thereafter processing the requested fragmentation size change to
7 modify a fragmentation size cluster communication parameter used in
8 transmitting messages from the source node to the plurality of target nodes;
9 and
10 (c) thereafter sending messages from the source node to the plurality
11 of target nodes using the modified fragmentation size cluster communication
12 parameter.
 - 1 2. The method of claim 1, further comprising sending a sync message from
2 the source node to the plurality of target nodes, wherein deferring processing of the
3 requested fragmentation size change includes waiting for an acknowledgment
4 message for the sync message from each of the plurality of target nodes.
 - 1 3. The method of claim 2, wherein the sync message is configured to initiate,
2 upon receipt by each target node, an immediate acknowledgment message from such
3 target node that acknowledges receipt for each unacknowledged message received by
4 such target node.
 - 1 4. The method of claim 2, wherein processing the requested fragmentation
2 size change includes sending a fragmentation size change message from the source
3 node to the plurality of target nodes, the fragmentation size change message
4 configured to modify the fragmentation size cluster communication parameter on each
5 of the plurality of target nodes.

1 5. The method of claim 4, wherein processing the requested fragmentation
2 size change further includes modifying the fragmentation size cluster communication
3 parameter on the source node.

1 6. The method of claim 4, wherein the source node comprises a message
2 queue, and wherein sending the sync message includes placing the sync message on
3 the message queue and sending the fragmentation size change message includes
4 placing the fragmentation size change message on the message queue.

1 7. The method of claim 6, wherein the source node and the plurality of target
2 nodes are members of a first group, wherein the source node is a member of a second
3 group that further includes as members a second plurality of target nodes, wherein the
4 source node further comprises a second message queue, wherein the source node
5 includes first and second connections respectively associated with the first and second
6 groups, wherein sending the sync message further includes placing the sync message
7 on the second message queue to initiate, upon receipt by each target node in the
8 second group, an immediate acknowledgment message from such target node that
9 acknowledges receipt for each unacknowledged message received by such target
10 node, and wherein sending the fragmentation size change message further includes
11 placing the fragmentation size change message on the second message queue to
12 modify the fragmentation size cluster communication parameter on each of the
13 plurality of target nodes in the second group.

1 8. The method of claim 7, further comprising sequentially issuing messages
2 on the first and second message queues respectively using first and second threads
3 executing on the source node.

1 9. An apparatus, comprising:

2 (a) a memory; and

3 (b) a program resident in the memory, the program configured to
4 dynamically modify a fragmentation size cluster communication parameter in
5 a clustered computer system by processing a requested fragmentation size
6 change only after receipt of an acknowledgment message for at least one
7 unacknowledged message sent by a source node to a plurality of target nodes.

1 10. The apparatus of claim 9, wherein the program is further configured to
2 process the requested fragmentation size change after receipt of the acknowledgment
3 message to modify a fragmentation size cluster communication parameter used in
4 transmitting messages from the source node to the plurality of target nodes, and
5 wherein the program is further configured to thereafter send messages from the source
6 node to the plurality of target nodes using the modified fragmentation size cluster
7 communication parameter.

1 11. The apparatus of claim 10, wherein the program is further configured to
2 send a sync message from the source node to the plurality of target nodes such that
3 deferring processing of the requested fragmentation size change includes waiting for
4 an acknowledgment message for the sync message from each of the plurality of target
5 nodes.

1 12. The apparatus of claim 11, wherein the sync message is configured to
2 initiate, upon receipt by each target node, an immediate acknowledgment message
3 from such target node that acknowledges receipt for each unacknowledged message
4 received by such target node.

1 13. The apparatus of claim 11, wherein the program is configured to process
2 the requested fragmentation size change by sending a fragmentation size change
3 message from the source node to the plurality of target nodes, the fragmentation size
4 change message configured to modify the fragmentation size cluster communication
5 parameter on each of the plurality of target nodes.

1 14. The apparatus of claim 13, wherein the program is further configured to
2 process the requested fragmentation size change by modifying the fragmentation size
3 cluster communication parameter on the source node.

1 15. The apparatus of claim 13, further comprising a message queue resident
2 on the source node, wherein the program is configured to send the sync message and
3 the fragmentation size change message by placing the sync message and the
4 fragmentation size change message on the message queue.

1 16. The apparatus of claim 15, wherein the source node and the plurality of
2 target nodes are members of a first group, wherein the source node is a member of a
3 second group that further includes as members a second plurality of target nodes, the
4 apparatus further comprising a second message queue resident on the source node,
5 wherein the source node includes first and second connections respectively associated
6 with the first and second groups, wherein the program is further configured to send
7 the sync message by placing the sync message on the second message queue to
8 initiate, upon receipt by each target node in the second group, an immediate
9 acknowledgment message from such target node that acknowledges receipt for each
10 unacknowledged message received by such target node, and wherein the program is
11 further configured to send the fragmentation size change message by placing the
12 fragmentation size change message on the second message queue to modify the
13 fragmentation size cluster communication parameter on each of the plurality of target
14 nodes in the second group.

1 17. The apparatus of claim 16, wherein the program is further configured to
2 sequentially issue messages on the first and second message queues respectively using
3 first and second threads executing on the source node.

1 18. A clustered computer system, comprising:

2 (a) a plurality of nodes coupled to one another over a network, the
3 plurality of nodes including a source node and a plurality of target nodes;

4 (b) a source program resident on the source node, the source program
5 configured to dynamically modify a fragmentation size cluster communication
6 parameter in a clustered computer system by sending a sync message to the
7 plurality of target nodes, thereafter waiting for an acknowledgment message
8 for the sync message from each of the plurality of target nodes, thereafter
9 sending a fragmentation size change message to each of the plurality of target
10 nodes to modify a fragmentation size cluster communication parameter on
11 such target nodes used in transmitting messages from the source node to the
12 target nodes; and

13 (c) a target program resident on each of the plurality of target nodes,
14 the target program configured to send an acknowledgment message to the
15 source node in response to the sync message to acknowledge receipt for each
16 unacknowledged message received thereby, and to modify a fragmentation
17 size cluster communication parameter associated therewith in response to the
18 fragmentation size change message.

- 1 19. A program product, comprising:
- 2 (a) a program configured to dynamically modify a fragmentation size
- 3 cluster communication parameter in a clustered computer system by
- 4 processing a requested fragmentation size change only after receipt of an
- 5 acknowledgment message for at least one unacknowledged message sent by a
- 6 source node to a plurality of target nodes; and
- 7 (b) a signal bearing medium bearing the program.

- 1 20. The program product of claim 19, wherein the signal bearing medium
- 2 includes at least one of a transmission medium and a recordable medium.

000001 00510960